

Use Ceph RBD Image as XenServer Local Storage Repository

Version: 0.9

Date: 20170627

I. Testing Environment:

1. XenServer 7.2. (build date 201705011, DBV 2017.0517)
Kernel version: Linux XenServer1 4.4.0+10
2. Mars200 Ceph Cluster, Ceph Jewel 10.2.7

II. Installation Steps:

1. Install XenServer 7.2
2. Install Ceph related package in XenServer
 - 2.1. ssh into XenServer,
 - 2.2. Install ceph-common package
 - Manually edit "/etc/yum.repos.d/CentOS-Base.repo" and enable Base package
 - run "yum repolist all"
 - run "yum install ceph-common"
 - 2.3. copy /etc/ceph/ceph.client.admin.keyring and ceph.conf to /etc/ceph directory
3. Check rbd image can be manually mapped in XenServer.
(solution1: direct map, 192.168.1.201 is ceph mon node's ip address, poolr2 is pool_name, image1 is image_name)

```
modprobe rbd
/bin/echo "192.168.1.201 name=admin,secret=AQB2XT5ZMRrJGxAAMWCNL04K3ekZDzG/hKbsaQ==
poolr2 image1" > /sys/bus/rbd/add
```


(solution2: use rbd map)

```
modprobe rbd
rbd map poolr2/image1
```
4. Set auto map/unmap rbd images on boot/shutdown.
 - 4.1. Edit /etc/ceph/rbdmap and add the pool/image you want to auto map

```
vi /etc/ceph/rbdmap
```


for example, add following setting into it

```
poolr2/image1 id=admin,keyring=/etc/ceph/ceph.client.admin.keyring
```
 - 4.2. Fix /etc/init.d/rbdmap error on XenServer 7.2

```
vi /etc/init.d/rbdmap
```


masked all log_action_begin_msg and log_action_end_msg lines (total 4 lines)
 - 4.3. load rbd module at startup

```
echo modprobe rbd >> /etc/rc.modules
chmod +x /etc/rc.modules
```
 - 4.4. enable rbdmap service

```
chkconfig
chkconfig rbdmap on
chkconfig
```
 - 4.5. restart the rbdmap service and check if image is correctly mapped as /dev/rbd

```
/etc/init.d/rbdmap restart
```
 - 4.6. reboot the server, and check if the image is correctly mapped on /dev/rbd device
5. Create Xen server Local Storage Repository using rbd image

- 5.1. `make /dev/rbd*` auto link to `/dev/disk/by-id`
`vi /usr/lib/udev/rules.d/50-rbd.rules`
#append following line in it
`KERNEL=="rbd[0-9]*", ENV{DEVTYPE}=="disk", PROGRAM="/usr/bin/ceph-rbdnamer %k",`
`SYMLINK+="disk/by-id/%c{1}-%c{2}"`
- 5.2. Re-start rbdmap service and check the symbol link is created in `/dev/disk/by-id`
`/etc/init.d/rbdmap restart`
`ls -l /dev/disk/by-id/`
- 5.3. Add rbd type into XenServer lvm config files.
 Note: XenServer add local disk default use lvm type. But the default lvm setting can't understand rbd type, so we need add rbd support to lvm configuration file.
 - a. Add into devices section of `/etc/lvm/lvm.conf`
`# types = ["fd", 16]`
`types = ["rbd", 64]`
 - b. Add into devices section of `/etc/lvm/master/lvm.conf`
`# types = ["fd", 16]`
`types = ["rbd", 1024]`
- 5.4. Create Local Storage Repository on XenServer master node (If you group XenServers in pools, do this in master node only)
`xe sr-create content-type=user device-config:device=/dev/disk/by-id/poolr2-image1`
`name-label="RBDSR" shared=true type=lvm`
- 5.5. Use "`xe pbd-list`" to check if pbd created or not.
`xe pbd-list`
- 5.6. Use XenCenter GUI to check if the storage is added correctly.
- 5.7. (Information) Delete Storage Repository
 - a. `xe sr-list name-label=<Name of the SR>`.
 - b. `xe pbd-list sr-uuid=<UUID of SR>`
 - c. `xe pbd-unplug uuid=<UUID of PBD>`
 - d. `xe sr-forget uuid=<UUID of SR>`

